Application of Grazing-Inspired Guidance Laws to Autonomous Information Gathering

Thomas Apker   Shih-Yuan Liu   Donald Sofge   J. Karl Hedrick

Abstract—**Domestic grazing animals follow simple, scalable rules to assign themselves trajectories to cover a pasture. We explain how to adapt these rules for an information gathering system based on a realistic robot motion model and Kalman-filter based evidence grid that accounts for both bandwidth and sensor limitations. Our results show that this algorithm can meet or exceed the performance of state of the art field robotics systems, particularly when scalability and robustness to failure are required.**

I. INTRODUCTION

Biological systems such as domestic grazing animals exhibit efficient and effective means of gathering food from the environment without explicit coordination or centralized direction [1]. This foraging process is fundamentally an area coverage task with many potential robotic applications ranging from surveillance to inspection. To have similar behavior in mobile robots, one must replace the communications channels used by the grazers and adapt the guidance rules to the kinetics of mobile robots.

There are two communications channels used to exchange the data required for grazing guidance, 1) peer locations and 2) food available. For ruminants, we assume for this work that the peer location channel comes from the animal's own senses and the physical location of its peers, and the "food available" channel comes from visual and olfactory measurements of local plant life. In other words, members of a herd or flock coordinate their grazing trajectories by changing positions and eating plants. We will present electronic replacements for these channels that allow the mobile agents access to the information they need to apply grazing-inspired algorithms.

Small mobile robots lack the ability to measure their peers locations accurately enough for coordination except at relatively short ranges using special visual cues [2], and so we assume that agents must report their positions over a shared network. We also assume that information gathering should not result in detectable changes to the environment.

Stable vehicle guidance requires vehicles being aware of their reachable space, leading to another difference between robotic grazing and its bio-inspiration. Animals appear to reason in a Euclidean space and subconsciously translate desired velocities into trajectories in their limbs' joint space, allowing humans to guide them with simple shepherding behaviors or remotely via electronic collars [3]. This allows a very broad range of accelerations and automatic adaptation to handle unreachable commands, allowing the animal's higher level reasoning to operate on positions and velocities instead of complicated, constrained kinetics. We will present steps that must be taken to allow grazing-inspired guidance algorithms to account for vehicle mobility limits.

To constrain the design space for this study, we will assume that the mobile robots each carry a downward-looking camera and that its images can be processed on-board using a classification algorithm to compress the data stream into relevant information. Motivated by weather-related disasters on both the East Coast of the United States and in the Philippines, we will assume that the robot herd has been deployed to evaluate the post-disaster state of a road network.

The remainder of this paper is divided into six sections. The next section describes the background and similar work. Section III describes how the changes in information exchange required for robotic applications are handled, including the relationship between network bandwidth and the types of information "food" the robot herd can collect. Section IV describes our approach to translate animal grazing trajectories into mobile robot guidance for area coverage. The performance of the algorithm is tested using a disaster recovery scenario in Section V and compared to flight paths designed to maximize coverage rate in Section VI. This is followed by a short conclusion in Section VII.

II. BACKGROUND

The most similar work to that presented here was conducted by Haumann *et al.* [4] on cooperative frontier exploration. In this work a small team of robots built up a map of an unknown environment by exchanging sensor data and computing control inputs that maximize their collection of information within a Voronoi decomposition of the space around each robot. While they define how to explicitly assign priority to particular regions for information gathering, they do not discuss the mechanics of fusing sensor data from multiple platforms.

Yamauchi [5] described an algorithm for multi-robot frontier exploration using low-fidelity ranging sensors, and the importance of designing the sensor and data fusion algorithms to avoid undesired behavior. His work assumed that the world can be represented by an information grid as

1. REPORT DATE **SEP 2014**	2. REPORT TYPE	3. DATES COVERED **00-00-2014 to 00-00-2014**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Application of Grazing-Inspired Guidance Laws to Autonomous Information Gathering	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Naval Research Laboratory ,Washington,DC,20375**	8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE Press, Chicago, IL, pp. 3828-3833, 14-18 Sep 2014.

14. ABSTRACT

Domestic grazing animals follow simple, scalable rules to assign themselves trajectories to cover a pasture. We explain how to adapt these rules for an information gathering system based on a realistic robot motion model and Kalmanfilter based evidence grid that accounts for both bandwidth and sensor limitations. Our results show that this algorithm can meet or exceed the performance of state of the art field robotics systems, particularly when scalability and robustness to failure are required.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a REPORT **unclassified**	b ABSTRACT **unclassified**	c THIS PAGE **unclassified**	**Same as Report (SAR)**	**6**	

described by Thrun *et al.* [6], broken into discrete cells with a finite probability of containing obstacles.

Julian *et al.* [7] present an information theoretic optimal routing algorithm for a team of mobile agents collecting information on such a grid assuming that each carries a sensor that can be expressed as Boolean detector of the signal of interest in each cell. This approach is too computationally intensive for many real-time applications, particularly on dense grids and with multiple vehicles, although Schwager [8] provides a grid-refinement method to help limit the computational cost. Kreucher *et al.* [9] used a combination of the information gradient and physics-inspired repulsion between agents to overcome this scalability constraint.

An alternative to the odds-based evidence grid is one that employs a Kalman filter or similar estimator that explicitly separates the uncertainty, generally expressed in terms of covariance, and the estimated value of the cell. Given a map of cell uncertainty, Yilmaz *et al.* [10] demonstrated a multi-vehicle approach to finding optimal paths by expressing it as the Selective Traveling Salesman Problem subject to dynamic constraints. Tisdale *et al.* [11] utilized a receding horizon approach along with sequential planning to generate optimized trajectories for the agents to maximize the information gain of the team.

In their study of grazing behavior, Liu and Hedrick [1] demonstrated that it is possible to replicate animal behavior using velocity commands derived from the status of nearby food and peer locations. Similar work by Balch *et al.* [12] on ants and Haque *et al.* [13] for dolphins and lions demonstrates similar results that suggest simple kinematic models are sufficient to explain animal behaviors.

However, rules that result in desirable behaviors for first order, kinematic systems do not in general converge to good solutions for second order, kinetic systems. Yu *et al.* [14] describe the conditions required to reach a consensus using a distributed controller in this context. The need for both reactivity and inclusion of kinetic constraints led the authors to choose the physicomimetics algorithm developed by Spears *et al.* [15] for this study.

III. COMMUNICATIONS AND DATA FUSION

Coordination of reactive foraging agents requires that they exchange data about 1) their own locations, and 2) information they have collected. If each agent can compute its own location using GPS or similar technology, then making it available to peers involves broadcasting a position once per decision cycle. This implies that each sensor information packet must contain the following information at its associated size s:

- Location, $s_l = 3 \times 32$bytes
- Orientation, $s_o = 32$bytes
- Uncertainty estimate, $s_P = 16 \times 32$bytes

A. Sensor Design

The remaining bandwidth, b_w, must be sufficient for N robots to transmit their sensor haul of s_H bytes once per decision cycle of length Δt_{outer}. At the extremes of the


Fig. 1. Plot of the packet size available to robotic foraging teams

sensor haul size are high definition (HD) video, which we will assume requires 2Mb/s for an $s_H = .2$Mb, and a Boolean disk per [9], with $s_H = 1$. The maximum size of this disc, d, was determined by the combination of the vehicle's altitude, h, and camera field of view α. The minimum packet size including the sensor haul and vehicle position information is shown in (1), and leads to team-wide bandwidth required b_w computed per (2). HD video is only possible with a high bandwidth network such as the IEEE 802.11g wireless Ethernet protocol, the cyan line in Fig. 1; while the Boolean disk sensor, red line in Fig. 1, allows a large team using longer range protocols such as the 900MHz xBee Pro.

$$s = s_H + s_l + s_o + s_P \quad (1)$$
$$b_w = Ns/\Delta t_{outer} \quad (2)$$

In between these two extremes are approaches that segment the image into discrete segments as shown in Fig. 2. In order to resolve features in the environment smaller than the sensor's diameter d, the sensor's resolution, d_{min}, must be at most half the size of the target signal; the minimum depth is a function of the vehicle's speed u and the sampling time Δt_{outer}. The impact of using a Boolean disk vs. a Boolean grid sensor is shown in Fig. 3; note that the false-positive rate of the Boolean disc's evidence grid is much higher than the Boolean grid result.

For this study, we will assume that some overlap during a single pass is desirable, and thus assume that the sensor is a square, 11×11 grid. When the content of that grid is a probability or percentage of occupancy, represented by a floating point value, then $s_H = 968$bytes; the bandwidth requirement follows the magenta curve in Fig. 1. We found that the best mix between sensor fusion performance and team scalability came with a Boolean grid, where each cell returned a binary value representing the presence of a target. In this case $s_H = 121$bytes, and the bandwidth required is represented by the blue curve in Fig. 1.

B. Data Fusion

The foraging models in [1] and [4] assumed that the agent will collect a unique sample from each point in the environ-

(a) Top View   (b) Front View

Fig. 2. Schematic of downward-looking camera image segmented into discrete regions.

Fig. 4. Plot of f_c of a log-odds cell given Boolean measurements.

(a) Log-odds Grid   (b) Kalman Filter Grid

Fig. 3. Map generated by a Boolean disc sensor on a log-odds grid (left) and Boolean grid sensor on a Kalman filtered grid (right).

ment which implies high confidence in the measurement. However, the most interesting applications for this approach involve inexpensive transducers and low-computation classifiers in a distributed network. Thus, our approach uses an information metric to define the amount of food available.

There are two basic options to convert the sensor data from the camera on an agent into information: odds-based and filter based. Odds-based techniques like [6] work well with Boolean disc sensors, but tend to produce low-resolution maps as shown in Fig. 3a. The information-as-food metric, f_c, is defined for each cell in (3), where $p(occ)$ is the probability of a cell being occupied. With this definition, the sensors are attracted to the nearest uncertain cell. With a log-odds grid, the cell value $\hat{x} \in (-\infty, \infty)$ must be rescaled to $p(occ)$ using (4) which requires a computation of the exponential function once per cycle per cell.

$$f_c = 1 - 2\|p(occ) - .5\| \qquad (3)$$
$$p(occ) = e^{\hat{x}}/(1 + e^{\hat{x}}) \qquad (4)$$

The main drawback of this approach is that cells do not always converge to a single value. Given contradictory evidence, $p(occ)$ may gain uncertainty and return to a value near 0.5 after re-sampling. In other words "eating" the information by sampling might *add* food to the system instead of removing it. Frontier-based methods avoid this problem by defining frontiers over many cells, but the

"greedy food" algorithm of [1] and the present work is unstable if contradictory measurements are possible.

Consider the case where a field cell produces a "True" measurement 50% of the time due to partial occupancy or a problematic sensor. Assuming that the sensor has a probability of detection p_d and a false alarm rate of p_f, and the k^{th} measurement, z_k, is given a value of one for positive observation and zero otherwise; the impact of each measurement on a cell can be expressed as (5). Using (3) to convert this into a food value and simulating sensor inputs allows us to predict the f_c in this hypothetical field cell in Fig. 4.

$$\hat{x}_{k+1} = \hat{x}_k + \begin{cases} \log \frac{p_d}{p_f} & \text{if } z_k = 1 \\ \log \left(\frac{1-p_d}{1-p_f} \right) & \text{if } z_k = 0 \end{cases} \qquad (5)$$

A more accurate model of "eating" information than an odds-based information grid is one that computes an estimate of the percentage of occupancy of a cell using a linear estimation process such as a Kalman filter. This has two distinct advantages over the odds-based version: 1) it is computationally faster since it does not involve rescaling via an approximate exponential function and 2) it explicitly separates the estimate from the quantity of information available. Since our target of interest involved a static signal source, we only considered the update portion of the filter. Furthermore, we can assume that each cell may be computed independently using (6) – (10). The variables used and their values as applied to this study are provided in Table I.

Symbol	Description	Value
y	innovation	*variable*
H	measurement model	1.0
S	covariance innovation	*variable*
P	covariance	*variable*
R	Measurement noise model	.8
K	Kalman gain	*variable*

TABLE I

VARIABLES AND VALUES FOR THE KALMAN FILTER CELL MODEL

Fig. 5. Plot of $f_c = P$ of a Kalman-filtered cell with Boolean updates

$$y = z_k - H\hat{x} \qquad (6)$$
$$S = HP_kH + R \qquad (7)$$
$$K = P_kHS^{-1} \qquad (8)$$
$$\hat{x}_{k+1} = x_k + Ky \qquad (9)$$
$$P_{k+1} = (1 - KH)P_k \qquad (10)$$

Using this approach, we can define the covariance of a field cell as its available food. This lets us guarantee monotonic reduction in food available with successive samples and control the rate at which the food "regrows," if at all, via a defined process model. Note that this approach computes the amount and/or probability of a target's cell occupancy, rather than the odds of a cell being occupied, and converges to approximately the expected value even with Boolean measurements, *i.e.* z_k equal to one or zero, in Fig. 5.

IV. MOBILITY

Fig. 6. Block diagram showing the physicomimetics algorithm's role in providing guidance to an autonomous vehicle.

To implement the movement commands generated by the grazing algorithm, we assumed a motion constrained second order agent model and computed desired velocities from the field. There are two components to consider: 1) the agent model and 2) the field model. The goal of the agent model is to represent the inner loop in Fig. 6 in the physicomimetics calculations in such a way as to minimize the error term $\vec{e}$. The field model translates the instantaneous velocity commands generated by the grazing algorithm into forces and torques that will drive the vehicle towards "food" without interfering with its neighbors.


Fig. 7. The guidance law

A. Agent Modeling

We based our agent model on a quadrotor helicopter such as an ARDrone2, which can be approximated as a holonomic particle described in [15] if, as in our case, its orientation does not affect performance. Its maximum speed was set to 1 m/s to limit off-design aerodynamic effects and allow moderate wind penetration at the commanded speed.

B. Field Modeling

1) Guiding A Quadrotor Agent: In this section the guidance law that guides an quasi-holonomic agent to a specific position on the field is explained in detail. As shown in Fig. 7, the position of the agent is denoted by $x \in \mathbb{R}^2$, and the nearest cell with food to the agent is located at $x_d \in \mathbb{R}^2$. The desired velocity of the agent is denoted by V_d and is defined to be proportional to the error in position with a constant factor of $\alpha > 0$, namely

$$V_d = \alpha(x_d - x). \qquad (11)$$

The virtual force applied to the agent is denoted by u and is defined to be proportional to the difference between the current velocity of the vehicle, $\dot{x}$, and the desired speed, V_d, with a factor of $k > 0$. More specifically,

$$u = k(V_d - \dot{x}). \qquad (12)$$

Using a point-mass model for the agent with virtual mass m, the input force u directly affects the 2nd derivative of the position of the agent according to

$$m\ddot{x} = u. \qquad (13)$$

Without lost of generality, we can assume that $m = 1$ and $x_d = 0$. By substituting (11) and (12) into (13), the dynamics of the closed-loop system can be written as:

$$\ddot{x} = -k\alpha x - k\dot{x}. \qquad (14)$$

By defining

$$x_1 = x \qquad (15a)$$
$$x_2 = \dot{x}, \qquad (15b)$$

the system can be written in state-space form as

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -k\alpha & -k \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}. \qquad (16)$$

The eigenvalues of this linear 2nd-order autonomous system are:

$$\lambda = \frac{-k \pm \sqrt{k^2 - 4k\alpha}}{2}. \qquad (17)$$

(a) Initial Condition   (b) Lawnmower Paths   (c) Grazing Pattern

Fig. 8.   Initial conditions, lawnmower paths and a grazing pattern from the simulation. The agent colors in Fig. 8b denote their differentiated roles.

Since $k > 0$ and $\alpha > 0$, the real parts of the eigenvalues are always negative. This shows that the guidance law results in a stable closed-loop system for all holonomic agents.

2) De-conflicting among Agents: Collision avoidance and area allocation are achieved using a Voronoi partition of the space with each agent being the generator of its Voronoi cell, as in [1] and shown in Fig. 8c. To improve the coverage rate, we introduced an additional heuristic: When an agent consumes all of the available information in its Voronoi cell, its V_d commands it towards the nearest food on the field.

V. Simulation and Testing

We implemented this algorithm in the MASON [16] simulation environment on a 960m×960m grid scaled to one cell per meter. The search zone is 500m×800m and centered in the field as shown in Fig. 8. The "roads" are the black tiles and have a minimum width of 11m. Prior to each episode's start we simulate "damage" by reassigning randomly selected blocks of 16×16 to be non-roads. The search zone is covered with an information grid with cells 6m on each side whose initial value is set to .5 and initial covariance to 1.0, both unit-less ratios.

The agents began in a north-south line centered on their "base" location of 150m×500m with enough fuel for 700s of flight time, allowing them to easily complete the longest lawnmower path in Fig. 8. The sensing model is a downward looking camera, collecting information from a square patch $d = 50m$ on a side with $d_{min} = 5m$ to resolve the 11m wide roads. The agent models were advanced in time using first order forward Euler integration. We simulated vehicle failures by halving the initial fuel level on selected agents. Note that simulations start with agents 50s travel time away from where they can begin collecting data.

VI. Results and Discussion

For this study, we were interested in three metrics, coverage, true positive rate (TPR), and false positive rate (FPR). The coverage is measured by aggregated covariance defined as $1 - \Sigma P(t)/\Sigma P(t = 0)$. A field cell was counted as a "positive" if its estimate exceeded .1, and a true positive if it contained one or more road tiles or a false positive if not. The inverse logic is true for negative cells. These metrics define the completeness and utility of the resulting roadmap.

The performance of the information gathering teams was measured and compared over one hundred simulation runs in the following three cases:
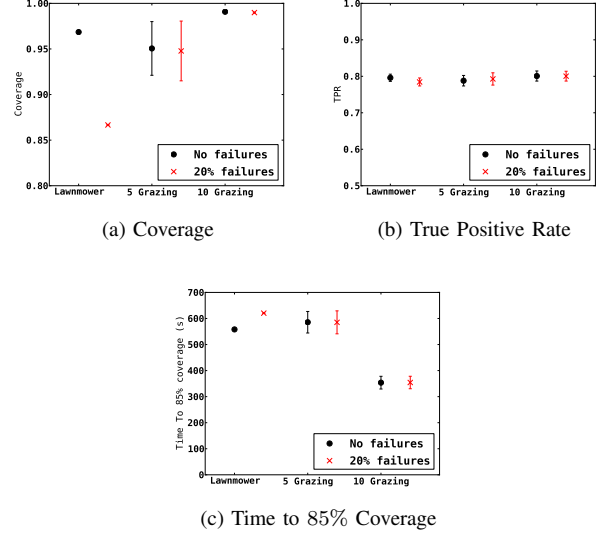

(a) Coverage

(b) True Positive Rate

(c) Time to 85% Coverage

Fig. 9.   Mean Coverage, True Positive Rate Performance, and Time to 85% of Change of Different Methods. Error bars indicate one standard deviation.

1) Five agents following lawnmower tracks with no failure and with 20% agent failures,
2) five agents grazing with no failure and with 20% failures,
3) and ten agents grazing with no failures and with 20% failures.

The coverage and TPR performance are shown in Fig. 9a and Fig. 9b respectively. The FPR is near zero in all cases and hence omitted.

In the lawnmower approach, each agent was attracted to a point 100m ahead of itself along its assigned track using the waypoint following logic in MOOS-IVP Helm [17]. The dynamics of the agents resulted in some oscillations about the track, but the team was still able to consistently achieve a coverage of 97% when no vehicles failed. Figure 10c shows that the best and worst performance at each timestep were nearly identical. When the agent with the longest track suffered the failure described above, coverage was reduced to 86% as shown in Fig. 9a. Note that the data quality metrics, specifically the TPR in Fig. 9b, were statistically identical in all cases studied.

A team of five of our grazing agents was able to "consume" between 84% and 98% of the covariance depending on the run, with an average consumption of 95% in both the failure and non-failure cases as shown in Fig. 9a. The substantial variance between runs reflects the impact of MASON randomizing the order of agent progation with each timestep. Occasionally, this would result in a cluster of agents such as the three near the top of Fig. 8c, where one agent would have an empty cell while chasing its peers towards the nearest food. The net effect on the team's performance is similar to an agent failure.

The real power of this algorithm is its scalability. As shown in Fig. 9a, a team that begins with ten grazing agents will consistently achieve 99% coverage. When the failure is

(a) Five Grazing Agents   (b) Ten Grazing Agents
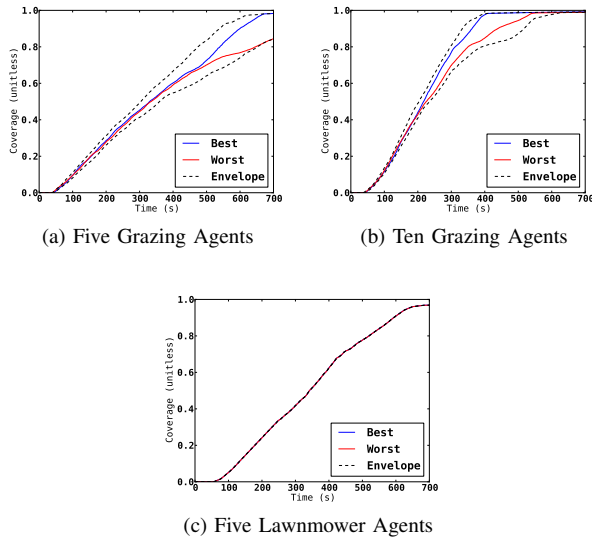

(c) Five Lawnmower Agents

Fig. 10. Coverage vs. time for the no failure cases. The Envelope curves show the range of coverage at each time-step over the 100 runs, while the Best and Worst curves show the runs resulting in most and least coverage, respectively.

applied at 350s, the team had covered 85% of the field, as shown in Fig. 9c, and spent the remaining half of the scenario pursuing isolated pockets of information. With more agents, the loss of one or two members to either a mechanical failure or unfavorable formation had less impact on the grazing performance.

Figure 10 shows the coverage vs. time for the cases with no failures. Note that Figs. 10a and 10b begin with a very tight envelope as the team spreads like a wave front until approximately 200s, when one or more agents may have exhausted its Voronoi cell. Since the "wavefront" with ten agents is much wider, the impact is less severe and occurs when the agents have already consumed nearly half the field. Also note that the Best and Worst curves do not lie on the Envelope until close to the end, as the team's coverage rate after the initial wave is a function of how widely scattered the "food" was and thus the amount of time agents spent in transit instead of consumption.

VII. Conclusion

When combined with a stable data fusion algorithm, our grazing-inspired collaborative area coverage performed almost as well as the state of the art in field robotic systems. Its reactive nature ensures complete coverage given sufficient time or additional agents. Imperfect path following due to agent dynamics does not require a substantial re-planning of trajectories to capture anything that was missed. Our proposed robotic grazing algorithm scales well, allowing large teams to accomplish information gathering tasks more quickly.

VIII. Acknowledgment

References

[1] S.-Y. Liu and K. Hedrick, "The application of domain of danger in autonomous agent team and its effect on exploration efficiency," in *American Control Conference (ACC), 2011*, 2011, pp. 4111–4116.

[2] E. Olson, "AprilTag: A robust and flexible visual fiducial system," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, May 2011, pp. 3400–3407.

[3] T. Wark, C. Crossman, W. Hu, Y. Guo, P. Valencia, P. Sikka, P. Corke, C. Lee, J. Henshall, K. Prayaga, J. 'Grady, M. Reed, and A. Fisher, "The design and evaluation of a mobile sensor/actuator network for autonomous animal control," in *Information Processing in Sensor Networks, 2007. IPSN 2007. 6th International Symposium on*, 2007, pp. 206–215.

[4] A. Haumann, K. Listmann, and V. Willert, "Discoverage: A new paradigm for multi-robot exploration," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*, 2010, pp. 929–934.

[5] B. Yamauchi, "Frontier-based exploration using multiple robots," in *Autonomous Agents*, 1998.

[6] S. Thrun and D. Burgard, Wolfram andFox, *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press, 2005.

[7] B. J. Julian, M. Angermann, M. Schwager, and D. Rus, "A scalable information theoretic approach to distributed robot coordination," in *Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on*, 2011, pp. 5187–5194.

[8] M. Schwager, B. J. Julian, and D. Rus, "Optimal coverage for multiple hovering robots with downward facing cameras," in *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, 2009, pp. 3515–3522.

[9] C. Kreucher, A. Hero, K. Kastella, and M. Morelande, "An information-based approach to sensor management in large dynamic networks," *Proceedings of the IEEE*, vol. 95, no. 5, pp. 978–999, May 2007.

[10] N. K. Yilmaz, C. Evangelinos, P. F. J. Lermusiaux, and N. M. Patrikalakis, "Path planning of autonomous underwater vehicles for adaptive sampling using mixed integer linear programming," *IEEE JOURNAL OF OCEANIC ENGINEERING*, vol. 33, no. 4, pp. 522–537, October 2008.

[11] J. Tisdale, H. Durrant-Whyte, and J. K. Hedrick, "Path planning for cooperative sensing using unmanned vehicles," in *Proc. ASME Int. Mechanical Engineering Conf. Exposition, Seattle, Washington*, 2007.

[12] T. Balch, F. Dellaert, A. Feldman, A. Guillory, C. Isbell, Z. Khan, S. Pratt, A. Stein, and H. Wilde, "How multirobot systems research will accelerate our understanding of social animal behavior," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1445 –1463, july 2006.

[13] M. Haque, A. Rahmani, and M. Egerstedt, "Geometric foraging strategies in multi-agent systems based on biological models," in *Decision and Control (CDC), 2010 49th IEEE Conference on*, Dec 2010, pp. 6040–6045.

[14] W. Yu, G. Chen, M. Cao, and J. Kurths, "Second-order consensus for multiagent systems with directed topologies and nonlinear dynamics," *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, vol. 40, no. 3, pp. 881–891, 2010.

[15] W. M. Spears, D. F. Spears, J. C. Hamann, and R. Heil, "Distributed, physics-based control of swarms of vehicles," *Autonomous Robots*, vol. 17, no. 2-3, pp. 137–162, 2004.

[16] S. Luke, C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan, "Mason: A multi-agent simulation environment," *Simulation: Transactions of the society for Modeling and Simulation International*, vol. 82, no. 7, pp. 517–527, 2005.

[17] M. R. Benjamin, H. Schmidt, P. M. Newman, and J. J. Leonard, "Nested Autonomy for Unmanned Marine Vehicles with MOOS-IvP," *Journal of Field Robotics*, vol. 27, no. 6, pp. 834–875, November/December 2010.